

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 31-12-2014		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 14-Dec-2011 - 13-Dec-2012	
4. TITLE AND SUBTITLE Final Report: Workshop on Planning and Learning in Multi-Agent Environments			5a. CONTRACT NUMBER W911NF-12-1-0020		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 1D10BP		
6. AUTHORS Dana S. Nau			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Maryland - College Park 3112 Lee Building College Park, MD 20742 -5141			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 60382-NS-CF.1		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This report documents research performed under ARO grant W911NF1210020 for the period 14 December 2011 through 13 December 2012. During this period we organized a workshop entitled Planning and Learning in Multi-Agent Adversarial Environments. The workshop was held on 26 and 27 April, 2012, and included 26 participants from academic institutions, companies, and government agencies. The format included speakers, breakout groups, and breakout-group reports. Subsequent to the workshop, a report summarizing the workshop results was prepared and distributed to the participants.					
15. SUBJECT TERMS planning, learning, multi-agent systems, robotics, adversarial reasoning.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Dana Nau
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 301-405-2684

Report Title

Final Report: Workshop on Planning and Learning in Multi-Agent Environments

ABSTRACT

This report documents research performed under ARO grant W911NF1210020 for the period 14 December 2011 through 13 December 2012. During this period we organized a workshop entitled Planning and Learning in Multi-Agent Adversarial Environments. The workshop was held on 26 and 27 April, 2012, and included 26 participants from academic institutions, companies, and government agencies. The format included speakers, breakout groups, and breakout-group reports. Subsequent to the workshop, a report summarizing the workshop results was prepared and distributed to the participants.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Dana S. Nau	0.00	
S.K. Gupta	0.00	
FTE Equivalent:	0.00	
Total Number:	2	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Total Number:

Names of personnel receiving PHDs

<u>NAME</u>
Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Technology Transfer

Workshop Report:

Planning and Learning in Multi-Agent Adversarial Environments

Organizers: Dana S. Nau
 Dept. of Computer Science
 University of Maryland
 nau@cs.umd.edu

 S. K. Gupta
 Dept. of Mechanical Engineering
 University of Maryland
 skgupta@umd.edu

Sponsor: Army Research Office
 POC: Dr. Purush Iyer
 purush.iyer@us.army.mil

Workshop date: April 26-27, 2012
Location: University of Maryland

Contents

1	Introduction	3
2	Format	3
3	Issues Identified by the Breakout Groups	3
3.1	Fast and Accurate Reasoning about Performance Constraints	4
3.2	Closed Versus Open World	4
3.3	Fast Online Planning and Replanning	4
3.4	The Symbol Grounding Problem	4
3.5	Predicting the Behavior of Intelligent Adversaries	5
3.6	Learning	5
3.7	Communication in Multi-Agent Teams	5
3.8	Challenges for Distributed Planning	6
3.9	Communication between robots and humans	6
3.10	Benchmark Problems and Training Data	6
4	Summary and Conclusions	7
A	Workshop Agenda	8
B	Planning and learning on physical robots, Day 1	9
C	Planning and learning on physical robots, Day 2	15
D	How to deal with intelligent adversaries, Day 1	20
E	How to deal with intelligent adversaries, Day 2	24
F	Multi-Agency in uncertain environments, Day 1	28
G	Multi-Agency in uncertain environments, Day 2	31

1 Introduction

An increasingly important problem is how to plan courses of action in multi-agent environments. Each agent's actions may change the environment in ways that will impact the other agents, hence each agent may need to reason about how its plans will affect (and be affected by) the others. This problem presents challenges that go beyond the capabilities of current AI planning and robotic planning techniques—especially in environments where there are both human and robotic agents, both friends and adversaries, and uncertainty about physical conditions such as terrain, weather, etc.

In existing research on robot planning—e.g., planning for AUVs—such sources of uncertainty are often represented by introducing large amounts of uncertainty into the outcomes of the AUV's actions. Such an approach may be sufficient to deal with adversaries that use simple brute-force strategies—but for success against intelligent adversaries, it is essential to incorporate ways to learn about the environment and the adversary, to make plans that take into account the likely actions of both the adversary and one's own team members, and to take into account the complex real-world requirements imposed by robotic platforms.

The purpose of the workshop was to explore how to develop new and more effective techniques for learning about the environment and the adversary, and for incorporating these techniques into mission planning and execution.

2 Format

The workshop ran for $1\frac{1}{2}$ -days, on April 26 and 27. The agenda (see Appendix 1) included eight presentations on April 26, followed by three concurrent breakout group discussions on the afternoon of April 26, and three more breakout groups on the morning of April 27.

On the first day, there were breakout groups on each of the following topics.

1. Implementing planning and learning on physical robots
2. How to deal with intelligent adversaries?
3. Multi-agency in uncertain environments

On the second day there were breakout groups on the same three topics, but with different group members and different group leaders in order to provide a different perspective. The six group leaders each prepared a summary of his/her group's discussion. The summaries are attached as appendices.

3 Issues Identified by the Breakout Groups

Even though the breakout groups were nominally on different topics, several issues came up repeatedly in the group discussions. The following subsections summarize the most important of these issues.

3.1 Fast and Accurate Reasoning about Performance Constraints

During each mission for which a robot is used, its capabilities will vary as the physical environment changes, or as the robot moves from one environment to another. Furthermore, the robot's capabilities will tend to degrade over time, due to decreasing battery power and to damage to the robot (e.g., sandstorms in Iraq caused the robots used by US forces to become unusable in about a week). *An important problem is how to reason about the robots' performance constraints when those constraints change over time.*

To meet this challenge, it is important to extend planning and learning research to incorporate physical simulations to get accurate predictions of the effects of the robot's actions. In addition to being accurate, these physical simulations will need to run exceptionally quickly, because they will need to be run multiple times during the planning or learning process, and the planning or learning process will be subject to hard real-time constraints.

3.2 Closed Versus Open World

Although current planning and learning algorithms incorporate a number of techniques for reasoning about uncertainty, these techniques generally depend on a "closed world" assumption, i.e., an assumption that all of the possible effects of each action are known in advance. In principle (though not necessarily in practice, see below), this assumption would make it possible to preplan an entire conditional plan or policy in advance. In practical robot planning, the closed-world assumption usually does not hold: instead, there may be anomalous events or anomalous action outcomes that are not present in the current world model. *The occurrence of open-world anomalies necessitates fast online replanning, and sometimes necessitates online creation of new or revised goals for the planning or learning algorithm.*

3.3 Fast Online Planning and Replanning

Most work on AI planning has focused on *preplanning*, in which a complete plan or policy is generated offline in advance before the robot begins executing it. In principle, such an approach is desirable because it would enable the robot to have a preplanned policy or contingency plan that it can execute quickly, rather than using an online planning algorithm that may incur execution delays. But as a practical matter, preplanning a complete plan or policy often is infeasible or impossible, due to two problems. Preplanning requires a closed-world assumption (see Section 3.2) that is quite unlikely to hold in adversarial robotic settings. But even when the closed world assumption holds, it often is infeasible to preplan an entire policy or contingency plan, due to exponential time and memory requirements. *Research is needed on effective techniques for generating partial plans and doing online plan refinement or replanning while execution proceeds.*

3.4 The Symbol Grounding Problem

AI planning and learning algorithms reason about symbolic goals, states, and actions. In robotic applications, goals and states and actions are collections of real numbers (e.g., (x, y, z) coordi-

nates). This creates substantial problems for AI planning and learning systems, in reasoning about whether a symbolic goal has been achieved, what constitutes an action model, and so forth. ***Better ways are needed to map between the abstract symbols used in AI planning and learning algorithms, and the numeric values used in robotics.***

3.5 Predicting the Behavior of Intelligent Adversaries

In current robotics research, approaches for planning in the presence of an adversary tend to neglect the adversary's reasoning capability, attempting instead to model the adversary's behavior as an increased amount of uncertainty. But in an adversarial environment, small changes to a robot's actions may lead to large consequences because the adversary may respond in different ways. Hence, ***an important challenge for adversarial robotic planning is how to reason effectively about the adversary.***

For example, techniques are needed for translating the physical aspects of an interaction (see Section 3.1) into the numeric utility values needed for game-theoretic calculations. Furthermore, the game-theoretic techniques themselves will require significant enhancements. Game-theoretic solution concepts (e.g., Nash equilibria) are not always useful because of they depend on assumptions (e.g., common knowledge of rationality) that are not always appropriate for the kinds of environments that we are discussing. More research is needed on how to learn useful predictive models of adversaries and their objectives. These models will need to be capable of dealing with possible deception by the adversary—for example, in order to detect whether an agent is an adversary or not.

3.6 Learning

Previous sections have already mentioned several important issues involving learning. These include the need to reason about a robot's changing physical constraints (Section 3.1), the symbol-grounding problem (Section 3.4), the need to learn in an open world (Section 3.2), and the need to learn predictive models of adversarial behavior and objectives, especially in the presence of possible deception by the adversary (Section 3.5). There are ***two additional issues not mentioned earlier: how to deal with temporal uncertainty in events and in outcomes, and how to provide speed and scalability in the presence of real-time execution constraints.***

3.7 Communication in Multi-Agent Teams

In current practice in robotic systems, reasoning about multiple agents is generally handled manually, and uncertainty about these agents largely ignored.

In environments where communication is reliable and there is a relatively low frequency of exogenous events, central planning (in which a single planning system generates plans for all of the team members) has some clear advantages. But effective central planning becomes much more difficult if communication links are unreliable, because a central planner may not become aware of problems quickly enough to respond to them. In addition to difficulties in establishing

and maintaining communication links, there also can be problems in communicating *about* the uncertainty: an agent may not even know its own state, let alone the states of the other agents.

In environments where there is significant communication uncertainty, central planning becomes quite difficult, because a central planner will not be able to communicate its plans reliably to the team members, nor to get reliable feedback about the plan execution status. Consequently, *dynamically changing environment with unreliable multi-agent communication necessitate distributed planning.*

3.8 Challenges for Distributed Planning

In distributed planning, coordinating the team members becomes a much more challenging problem. Here are some examples:

- If an agent is damaged or destroyed, how can its assigned tasks be reassigned to the other agents?
- If an agent is still functional but is having difficulty carrying out its assigned task, when is it appropriate for the agent to break its commitment, and how can the task be reassigned to other agents?
- If a team receives additional agents (e.g., by reassignment from another team whose tasks have been finished), after they have begun performing their tasks, how should tasks be reallocated in the larger team?

3.9 Communication between robots and humans

Additional communication difficulties occur in teams that include both humans and robots. One is the difficulty of finding the right level of communication abstraction. For example, consider the task of flying a drone. Soldiers generally ask for situation assessments such as video feeds, that are raw data rather than interpreted communication; but the difficulty of assimilating this data means that a large number of humans (currently 12) are needed to fly the drone. Interesting events may be very few, but it would be desirable to have effective ways for robotic agents to perform other non-interesting things more autonomously.

3.10 Benchmark Problems and Training Data

For many of the research tasks outlined above, it will be important to have shared training data and benchmark problems. Such a collection of benchmark problems will need to balance to competing needs: the need to remove distracting technical details in order to carry out research tasks effectively, and the need for data and benchmarks that are realistic enough that the research results will have an impact on real-world robotics.

Real-time strategy games may provide useful data for research on predictive models of adversaries, modeling long-term and short-term plans, and incorporating the effects of forming and

shifting alliances. Tactical games such as the RoboCup competition may also be useful sources of data.

4 Summary and Conclusions

As discussed in the previous section, the workshop attendees identified a number of research issues related to planning and learning in multiagent adversarial environments. Here is a quick summary of the main points:

- An important yet neglected problem is how to reason about the robots performance constraints when those constraints change over time.
- The occurrence of open-world anomalies necessitates fast online replanning, and sometimes necessitates online creation of new or revised goals for the planning or learning algorithm.
- Research is needed on effective techniques for generating partial plans and doing online plan refinement or replanning while execution proceeds.
- Better ways are needed to map between the abstract symbols used in AI planning and learning algorithms, and the numeric values used in robotics.
- An important challenge for adversarial robotic planning is how to reason effectively about the adversary.
- Some important challenges for learning algorithms include how to deal with temporal uncertainty in events and in outcomes, and how to provide speed and scalability in the presence of real-time execution constraints.
- Dynamically changing environment with unreliable multi-agent communication necessitate distributed planning.
- In distributed planning, coordinating the team members becomes a much more challenging problem.
- Additional communication difficulties occur in multi-agent teams that include both humans and robots.
- It will be important to have shared training data and benchmark problems.

Thursday, April 26

0830	Registration/Continental Breakfast	
0900	Purush Iyer, ARO	<i>Welcome</i>
0915	Introduction of all participants	
0935	Dana Nau, UMD	<i>Game-Theoretic Planning in Partially Observable Euclidean Space</i>
1000	Robert Goldman, SIFT, LLC	<i>Planning Autonomous Agency: Reaction, Projection, and Hybrids</i>
1025	--- Break ---	
1040	Brad Clement, NASA Ames	<i>Multiagent Planning Space Applications</i>
1105	Chris Geyer, iRobot	<i>Challenges for Robotics in Adversarial Environments</i>
1130	Ed Durfee, U. Michigan	<i>Reasoning About Predictability in Cooperative and Adversarial Environments</i>
1155	--- Lunch in Room 2117 ---	
1245	Ashutosh Saxena, Cornell University	<i>Integrated Perception and Planning for Robots</i>
1310	Stuart Young, US Army	<i>Army Problems in Robotics</i>
1335	S.K. Gupta, UMD	<i>Learning Opportunities in Physics-Aware Planning</i>
1400	--- Break ---	
1415	Breakout group discussions*	
1545	Preparation of breakout group summaries	
1600	--- Break ---	
1620	Working group summaries (10 minutes each)	
1650	Discussion	
1710	--- End ---	

Friday, April 27

0830	Continental Breakfast	
0900	Purush Iyer, ARO	<i>Charter for the day</i>
0915	Breakout group discussions	
1045	--- Break ---	
1100	Working group summaries (20 minutes each)	
1200	--- Lunch ---	
1245	Discussion	
1300	--- Close Meeting ---	

B Planning and learning on physical robots, Day 1

Group leader: Robert Goldman

ARO workshop breakout: planning and learning in robots

Robert P. Goldman

2012-04-26 Thu

Shared training data

- ▶ Need training data
 - ▶ Challenge is to cover parts of the state space that the human doesn't encounter
E.g., if we drive off the road, then the robot won't know what the the agent *should* do when off the road.
 - ▶ Would need to have sensors needed for perception for autonomy
 - ▶ Need to information about the accuracy of the sensors
- ▶ Test cases (inside simulation)
 - ▶ High fidelity simulators that aren't too slow
 - ▶ Open source — practically you need access to the internals
 - ▶ Perception is the weak part of most currently available simulators. Poor noise models.
- ▶ Opportunity to actuate in the real world

Need to get stuff to run really, really fast in order to run on robots

- ▶ Planning at different granularity to be real-time
- ▶ Speedup learning problem
- ▶ Anytime behaviors
- ▶ Abstraction

Integrating a high-fidelity simulator with higher-level reasoning

- ▶ Directing sampling – find relevant parts of the space
 - ▶ Adaptive sampling
- ▶ Speed of simulation
- ▶ Abstraction
 - ▶ Learning abstraction
 - ▶ Learning models for abstractions

Symbol-grounding problem

- ▶ Classes
 - ▶ Perception
 - ▶ Action
 - Related to plan recognition/intent recognition

Learning of symbolic, projective action models

Relation of case-based and model-based planning

Failure modes of hierarchical systems

- ▶ Credit assignment across agent layers

Adversarial reasoning

- ▶ Categorize adversaries
- ▶ Knowing that adversary may be trying to deceive you
- ▶ The zero-day attack
 - ▶ “Non-parametric” defense (minimax vs. adversary model)
 - ▶ Explanation-based learning, ILP, one-instance learning
- ▶ Planning for deception

Exploiting multiple agents for learning

- ▶ Try a portfolio actions
 - Active learning
- ▶ Accept some loss
 - We can afford to lose some autonomous platforms to learn
- ▶ Not done in existing, e.g., RL techniques

Team versus team planning

- ▶ Detect team membership

Evolving model of platforms

C Planning and learning on physical robots, Day 2

Group leader: Chris Geyer

8 Themes

1. Grounding goals
2. Taking the initiative
3. League of benchmark problems
4. 1000's of agents – improving scalability
5. Balance replanning and preplanning
6. Conformant planning
7. Stealth, inconspicuity, and being non-threatening
8. Warrior LifeLog

Theme 1 – Grounding Goals / Goal Recognition

- #1 problem applying abstract approaches to real world robots is recognizing when abstract goals have been achieved
 - In AI community often assumed solved – functions evaluate when objective achieved – yes/no
 - In robotics community goals are xyz – coordinates
 - How do you avoid the long tail of point solutions
 - How do you generalize goal recognition
 - “wor won,” “enemy contained,” “enemy destroyed”

Theme 2 – Taking the initiative

- How you create systems that make goal or sub-goal proposals
- Given current state and conditions can a robot learn (from data or perhaps demonstration) what are an operator’s likely goals, and then make those proposals

Theme 3 – League of benchmark problems

- Goal: Develop a league of benchmark problems – with different technical foci – that can be used to do demonstrate approaches' efficacies while abstracting away distracting technical details
- Balance impact on real world robots vs. distracting technical details
- E.g., RoboCup has different leagues each of which focus on different technical (small & embedded; legged; humanoid, etc.)

Theme 4 – How to scale to 100s to 1000s of agents

- Currently approaches like MDP are limited to very low DoF problems
- Goal: Develop approaches that scale approaches to handle large numbers of robots – 100s to 1000s of robots
- E.g., approaches that learn policies on subsets of larger problems, and that can be combined in non-trivial way to one joint policy

Theme 5 – Balance replanning and preplanning

- Computationally limited systems or those faced with high DoF often precompute plans – however often not enough when faced with current conditions, necessitating replanning
- How do you balance replanning and preplanning?

Theme 6 – Conformant planning

- Conformant planning has been developed in symbolic planning community – can be adopted or adapted to real world robots?
- Adversaries will seek to narrow your options (make effects inevitable, use deception, etc.) – how do you plan so as to maintain your options?
- & make applicable to real systems

Theme 7 – Stealth, inconspicuity, and being non-threatening

- How do you make robots
 - Stealthy;
 - Inconspicuous; or,
 - Just non-threatening?
- In latter case, not just about not being detected but about not intimidating others, not inviting attack
- Conversely, how do you act aggressively?

Theme 8 – Warrior LifeLog

- We are interested in solving multi-agent adversarial
- Why not...
 - Put sensors (cameras, GPS, mics) on all warriors, vehicles and their weapons in Red Team/Blue Team exercises
 - Annotate it
 - Mine the data
 - Learn plans for multi-agent adversarial environments
 - How people communicate to achieve
 - Can you improve Army's operations – comms between people, TTPs, etc., etc.

D How to deal with intelligent adversaries, Day 1

Group leader: Mary Ann Fields

How to deal with intelligent adversaries

- Need to move beyond the “rational” opponent and Nash equilibria to the idea of a mixed environment with varying degree of rationality over opponents (robot and human adversaries) In the real world there is often only partial information and predicting adversarial patterns over time
- Look at games over a complexity scale that includes discrete turn based to continuous games, single and multiple opponents and differing levels of cooperation
- One important question is how do we predict adversarial behavior? How do we recognize adversarial behavior In some situations we are able to classify behavior in a small number of and respond to those
- In multi – player environments, we need to look at concepts of cooperation and trust – as ed pointed out not every agreement between players is broken deliberately - how does that effect the long term cooperation of entities
- Important question is how do you abstract plans from observation of low level actions.
- One tool to use to study conflicts is real time strategy games in which there are multiple opponents that may form temporary alliances, generally the environment is only partially observable. Right now the emphasis is not on building agents to win the game but to win vignettes from the game
- Finally we need to look at different domains in which to study intelligent adversaries –
- In strategic planning we might look at strategic computer games including predictive models of adversaries, modeling long term and short term plans, incorporating the effects of alliances, forming, shifting alliances
- In tactical games , martial arts , robocup, and sports might give us fruitful environments for studying adversaries

How to deal with intelligent adversaries

- Need to move beyond the “rational” opponent and Nash equilibria to the idea of a mixed environment with varying degree of rationality over opponents (robot and human adversaries) In the real world there is often only partial information and predicting adversarial patterns over time
- Look at games over a complexity scale that includes discrete turn based to continuous games, single and multiple opponents and differing levels of cooperation
- In multi – player environments, we need to look at concepts of cooperation and trust – asd pointed out not every agreement between players is broken deliberately - how does that effect the long term cooperation of entities
- Important question is how do you abstract plans from observation of low level actions.
- One tool to use to study conflicts is real time strategy games in which there are multiple opponents that may form temporary alliances, generally the environment is only partially observable. Right now the emphasis is not on building agents to win the game but to win vignettes from the game

Predicting/recognizing plans

- One important question is how do we predict adversarial behavior? How do we recognize adversarial behavior? In some situations we are able to classify behavior is a small number of and respond to those
 - Can we develop an “dead reckoning” algorithm for predicting adversary’s plans
-
- Learning to abstract plans from low level actions
 - How does various methods scale to multi-player environments

How to study the problem

- One tool to use to study conflicts is real time strategy games in which there are multiple opponents that may form temporary alliances, generally the environment is only partially observable. Right now the emphasis is not on building agents to win the game but to win vignettes from the game

Challenge Problems

- Multiple domains
- Strategic Games
 - incorporate effects of alliances,
 - forming, shifting alliances
 - What about unintentional actions that violate the alliance
 - Trust
 - Uncertainty
- Tactical game

E How to deal with intelligent adversaries, Day 2

Group leader: Ed Durfee

Characterizing the Adversary

- Static or stupid
 - Stochastic
 - Adaptive/Dynamic
 - Reflective/Recursive
 - Strategic/Long-term
-
- Hard to identify friend/foe/neutral
 - Adversary is not always adversarial

Objectives of the Adversary

- Disrupt our activities
 - Make the environment unsuited to doing what we want to do
 - Could be more tactical
- Achieve their own objectives
 - Sequential decisions (plans) with narrower intent
 - Could be more strategic

Robustness to Adversary

- Compiled “rules”
 - E.g., be unpredictable by default, minimize chokepoints
 - Danger: Lose reasons for behaviors, hard to quickly adapt when mismatch with situation faced
- First principles
 - Build/learn/maintain models of adversary and environment
 - Effects-based reasoning, accounting for adversary
 - Danger: Hard to populate models, non-stationary aspects of models, deception
 - Danger: Slow to use
- Minimax
 - Model adversary as worst-case
 - Danger: Overestimate enemy; high cost or infeasibility of achieving mission goals while staying “safe”

Responses to Adversaries

- Exploit technological superiority:
 - Defeat by utilizing lots of (robotic) assets
- Exploit tendencies of enemy
- “Train” the enemy

Technologies Available

- Game Theory
- Plan/Intent Recognition
- Sequential Decision Methods
 - Use current model to develop branching contingency plan with actions conditioned on (recursive) belief state (a policy)
 - Execute plan, collecting statistics on experienced transitions, opponent's behaviors, etc.
 - Update model and repeat
- Recursive agent modeling
- Machine Learning
 - Availability and representativeness of training examples...

Tools for User

- Decision-Support technologies:
 - Help manage multiple objectives under multiple threats
 - Detecting tendencies in opponent's behaviors
 - Detecting tendencies in own behaviors
- Robotic capabilities
 - Capabilities to get adversary to reveal information (e.g., smoke enemy out)

Convincing End Users

- Scope mission
 - Protecting an area?
- Scope adversary
 - Disruptive, but not deceptive?
- Evaluation Methodology
 - Simulated wargames with human opponents
 - Field exercises
 - Utilize military red-team experts

F Multi-Agency in uncertain environments, Day 1

Group leader: Brad Clement

SOA

- Typical handling of MA uncertainty is human/manual (as is most MA planning/learning).
- Deployments and more mature research tends to ignore uncertainty or replan.
- Research
 - slack
 - predictability
 - coupling
 - making/breaking commitments
 - trust

How to leverage research?

- Low-hanging fruit
- Bite off easier, more relevant problems
- Agents as sensors for reducing uncertainty

Challenges

- Designing agent systems to control uncertainty
 - How many agents?
 - Heavy duty or lightweight?
- Uncertainty of communication
- Modeling humans and how comm influences them
- Human-robotic interactions - autonomy
- Trust
- When to break commitments
- Who to talk to when receiving unexpected information
 - Passing troops say “don’t go there?”
 - Need of context/importance of order
- Processing massive data (collected in a MAS).

Approaches

- Focal points of coordination
 - On what to coordinate/synchronize/comm?
 - E.g., finding rendezvous points
 - Reasonable fall-back point
- Crowd sourcing
- DARPA Mind’s Eye
 - Learning/detecting MA actions

Scalability Issues

- Difficult to make decisions when simulating an uncertain future is intractable
- Abstraction of uncertainty

End users

- Need to know ramifications

G Multi-Agency in uncertain environments, Day 2

Group leader: Prasad Tadepalli

Uncertainty in Multiagent Environments

- Adversarial: small differences might lead to large consequences, because the opponent's might try to exploit them.
- Uncertainty about the opponent's intentions or policies
- Observations are small
- Explicit Deception
- Non-adversarial domains – small errors usually only lead to small differences

What uncertainties matter?

- Some uncertainties don't matter. Others do.
- Can we quantify where it is going to be high and where it does not matter?
- Not enough theory

On-line vs Offline Planning

- Most planning work is offline
- RL begins with no knowledge; does on-line greedy execution, but not much planning
- We need more model learning and on-line planning and correction of models
- Exploration vs exploitation. Learning factored models with uncertainty
- Must take calculated risks
- Energy efficiency – following is easy if you have a model; tracking is inefficient

Advantages of Multiagency

- In multiagent environments you can be more tolerant due to multiple agents on your team;
- Can use multiple agents to gather information
- *Reduce* uncertainty and *accommodate* uncertainty
- State space can be reduced by ignoring agents that are far away
- State evaluation mechanism can sometimes be accurate and sometimes not
- Evaluation function based on material is not good in unstable situations – search longer to take care of this problem

Temporal Degradation

- Factoring in risk is something you should do
- Robotics – performance degrades over time
- Uncertainty model may become obsolete
- What can you do with degraded components?
- How to update models? Must use a combination of online and offline models.
- Judicious mixture of offline and online. Robots performance degrades with battery

Robustness to changes

- Sandstorms in Iraq – made the robots unusable in a week
- Reflection of the bridge prevents
- Lighting conditions make the sensors not work.
- Add and remove new robots: Coverage planning studied, but not in general planning problems
- Worst-case vs, expected case – both are non-ideal
- With adversaries, it is more difficult

Deception

- Adversaries with deception is even more difficult.
- Complexity of the task itself might make it complicated
- Learning can help with lack of knowledge of pdfs but not deception. Some weaknesses can be exploited.
- Deception can only be beaten by deception
- Consider multiple hypotheses.
- Mixed strategy – a person could be deceiving at any time;
- Maintain an internal model of the opponent

Multi-agent uncertainty

- Multiagent coordination with uncertainty
- Predicting what other agents are trying to do
- Addressed communication, but not realistic communication about *uncertainty* and conditions of uncertainty
- Commitments in task distribution – agent A has difficulty in doing a task and B needs it, when A can violate his commitment or ask for help.
- Agent may not know its own state and its own capability let alone the others' states and plans.

Multi-agent uncertainty

- Agents know very little about each other and about their own state.
- Formation flying: need different kind of synchronization of knowledge of each others' state. Time synchronized. Even 100; ms difference can be a big problem.
- RL has not paid much attention to temporal uncertainty.
- Hurricane monitoring – everybody sampling in a different state. Even small lags generate bad models. Certain formations are good in preventing errors.

Decentralized vs Centralized Planning

- A lot of agents with centralized control – a lot of uncertainty about the agents is not there
- Works well if there is a reliable communication between robots.
- Centralized planning works in many cases but in DOD setting, we have to think about their security
- Mixed setting of robots and humans – how to communicate uncertainty between the humans and robots.
- Humans may be good at estimating uncertainty and lousy in predicting risk
- PDF-based uncertainty communication is difficult

Communication

- Soldiers need to be trained in communicating about uncertainty
 - How does the robot explain its conditions
 - One thing that soldiers ask for is situation assessment – give me a video feed \; mostly raw data – not interpreted communication
 - 12 people take to fly a drone for the airforce
 - Perception is the biggest weakness here
 - How to prevent risk in life-or-death situations?
 - Interesting events may be very few, but other non-interesting things may be done more autonomously
-
- Looking for needle in the hay-stack
 - Monitoring multiple screens. Information ground-up is important.